



(12) **United States Patent**
Boertjes

(10) **Patent No.:** **US 9,455,788 B2**
(45) **Date of Patent:** **Sep. 27, 2016**

(54) **HITLESS MODULATION SCHEME CHANGE SYSTEMS AND METHODS IN OPTICAL NETWORKS**

(71) Applicant: **David W. Boertjes**, Nepean (CA)

(72) Inventor: **David W. Boertjes**, Nepean (CA)

(73) Assignee: **Ciena Corporation**, Hanover, MD (US)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 83 days.

7,792,044 B2	9/2010	Holness et al.
7,826,752 B1	11/2010	Zanoni et al.
8,009,692 B2	8/2011	Basso et al.
8,009,985 B1	8/2011	Roberts et al.
8,295,175 B2	10/2012	Ellis et al.
8,364,036 B2	1/2013	Boertjes et al.
8,402,121 B2	3/2013	Skalecki et al.
8,467,687 B2	6/2013	Sasaki et al.
8,509,621 B2	8/2013	Boertjes et al.
8,553,707 B2	10/2013	Swinkels et al.
2004/0156644 A1	8/2004	Fuse et al.
2009/0214221 A1*	8/2009	Li et al. H04B 10/40 398/136
2011/0013911 A1	1/2011	Alexander et al.

(Continued)

(21) Appl. No.: **14/534,657**

(22) Filed: **Nov. 6, 2014**

(65) **Prior Publication Data**

US 2015/0229404 A1 Aug. 13, 2015

Related U.S. Application Data

(63) Continuation-in-part of application No. 14/176,908, filed on Feb. 10, 2014, now Pat. No. 9,258,190.

(51) **Int. Cl.**

H04B 10/516 (2013.01)
H04L 29/08 (2006.01)
H04L 12/24 (2006.01)
H04Q 11/00 (2006.01)

(52) **U.S. Cl.**

CPC **H04B 10/516** (2013.01); **H04L 69/321** (2013.01); **H04L 69/323** (2013.01); **H04L 69/324** (2013.01); **H04L 41/0896** (2013.01); **H04Q 2011/0086** (2013.01)

(58) **Field of Classification Search**

CPC combination set(s) only.
See application file for complete search history.

(56) **References Cited**

U.S. PATENT DOCUMENTS

6,459,832 B1 10/2002 Smith et al.
6,795,394 B1 9/2004 Swinkels et al.

OTHER PUBLICATIONS

Ciena Corporation, "Enabling a Bigger and Smarter Network With Ciena'S Wavelength 3", Application Note, AN085 8.2013.

(Continued)

Primary Examiner — David Payne


Assistant Examiner — Omar S Ismail

(74) *Attorney, Agent, or Firm* — Clements Bernard PLLC; Christopher L. Bernard; Lawrence A. Baratta, Jr.

(57) **ABSTRACT**

A hitless modulation change method at a node in an optical network includes determining that a modulation change is warranted for an optical modem in the node, the optical modem configured to communicate over an optical link; determining an impact of the modulation change on the optical link and associated underlying connections thereon; causing changes in a data plane for the associated underlying connections, prior to performing the modulation change; and causing the modulation change subsequent to accommodating the associated underlying connections in the data plane, thereby minimizing interruptions of the associated underlying connections due to the modulation change.

20 Claims, 8 Drawing Sheets



(56)

References Cited

U.S. PATENT DOCUMENTS

2011/0158642 A1* 6/2011 Barnard et al. H04B 10/572
398/25
2011/0222394 A1 9/2011 Swinkels et al.
2011/0222846 A1 9/2011 Boertjes et al.
2012/0163820 A1 6/2012 Dangui et al.
2012/0219288 A1 8/2012 Roberts et al.
2012/0219293 A1 8/2012 Boertjes et al.
2013/0209089 A1 8/2013 Harley et al.
2013/0209091 A1* 8/2013 Mateosky
et al. H04B 10/07953
398/26

2013/0216232 A1 8/2013 Zanoni et al.
2013/0236169 A1 9/2013 Gaudette et al.
2013/0272318 A1 10/2013 Swinkels et al.
2013/0308948 A1 11/2013 Swinkels et al.

OTHER PUBLICATIONS

Gho et al., "Rate-Adaptive Coding for Optical Fiber Transmission Systems," IEEE Journal of Lightwave Technology, vol. 29, No. 2, January 15, 2011.

* cited by examiner




FIG. 1


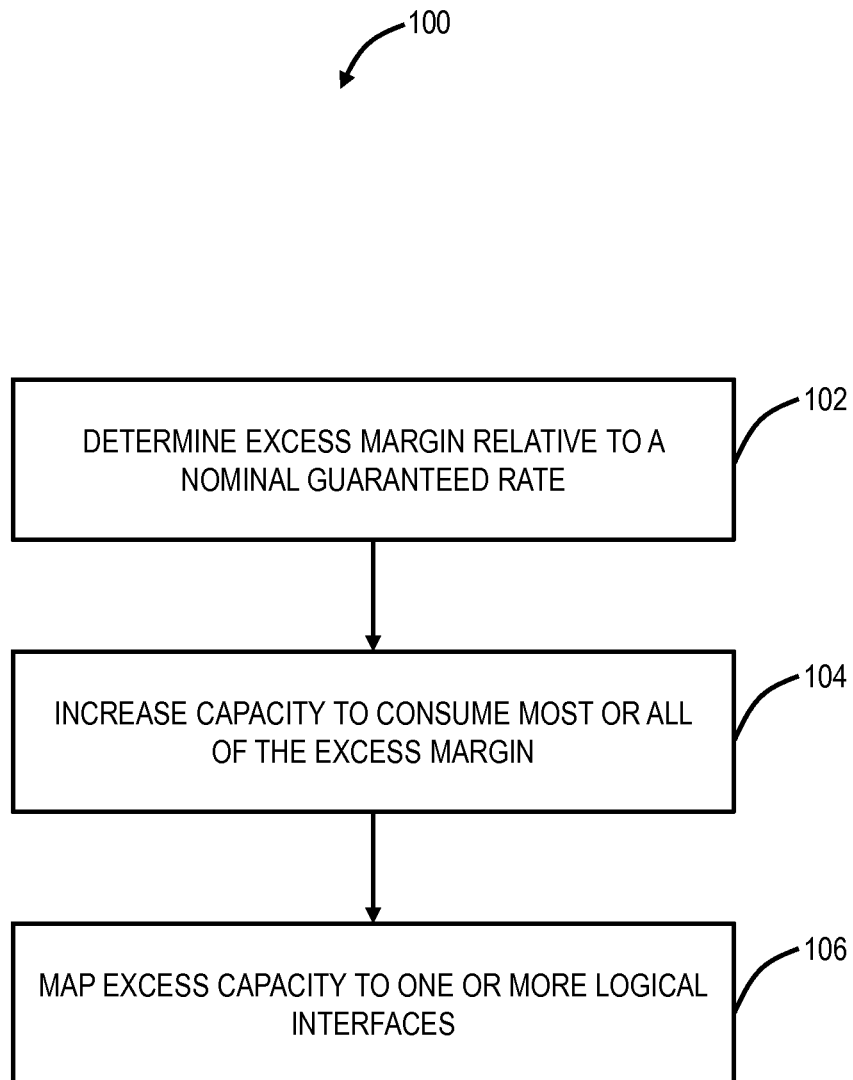


FIG. 2

**FIG. 3**

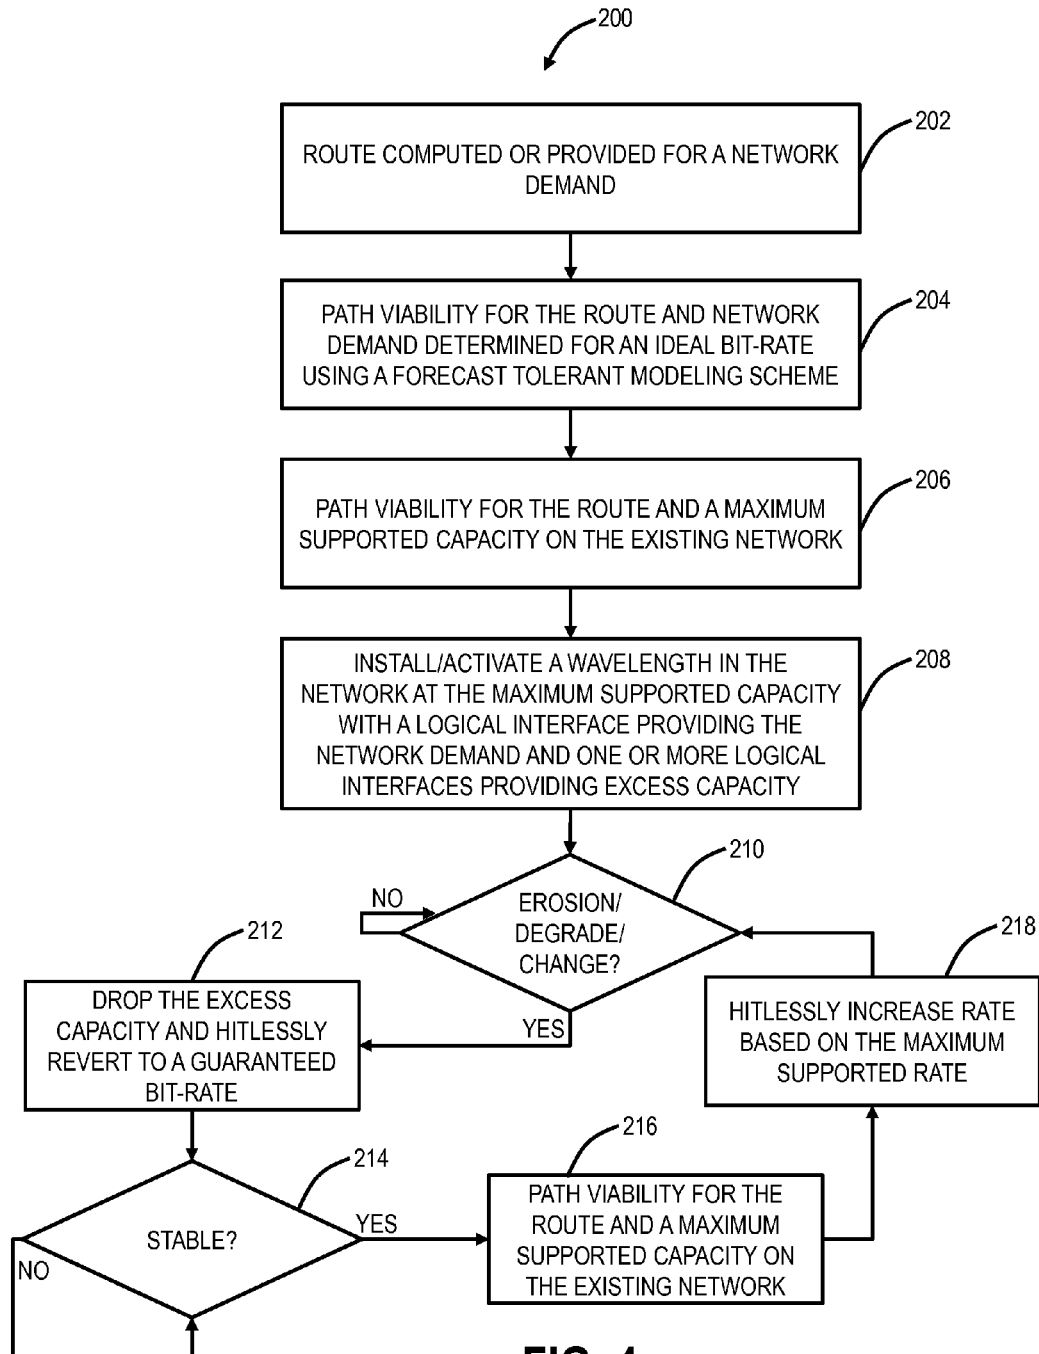
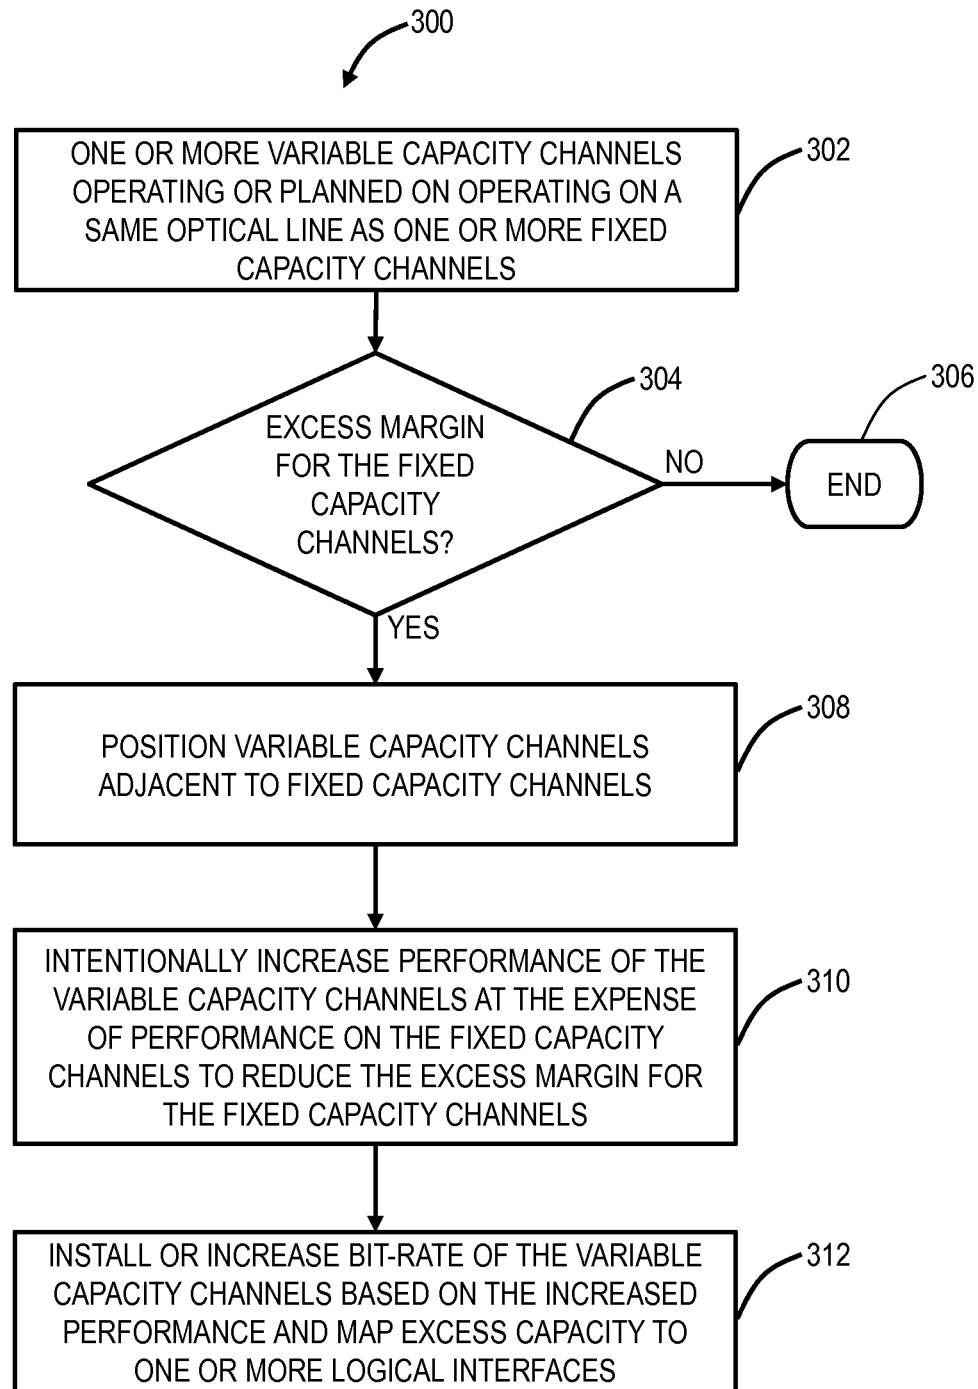


FIG. 4

**FIG. 5**




FIG. 6




FIG. 7


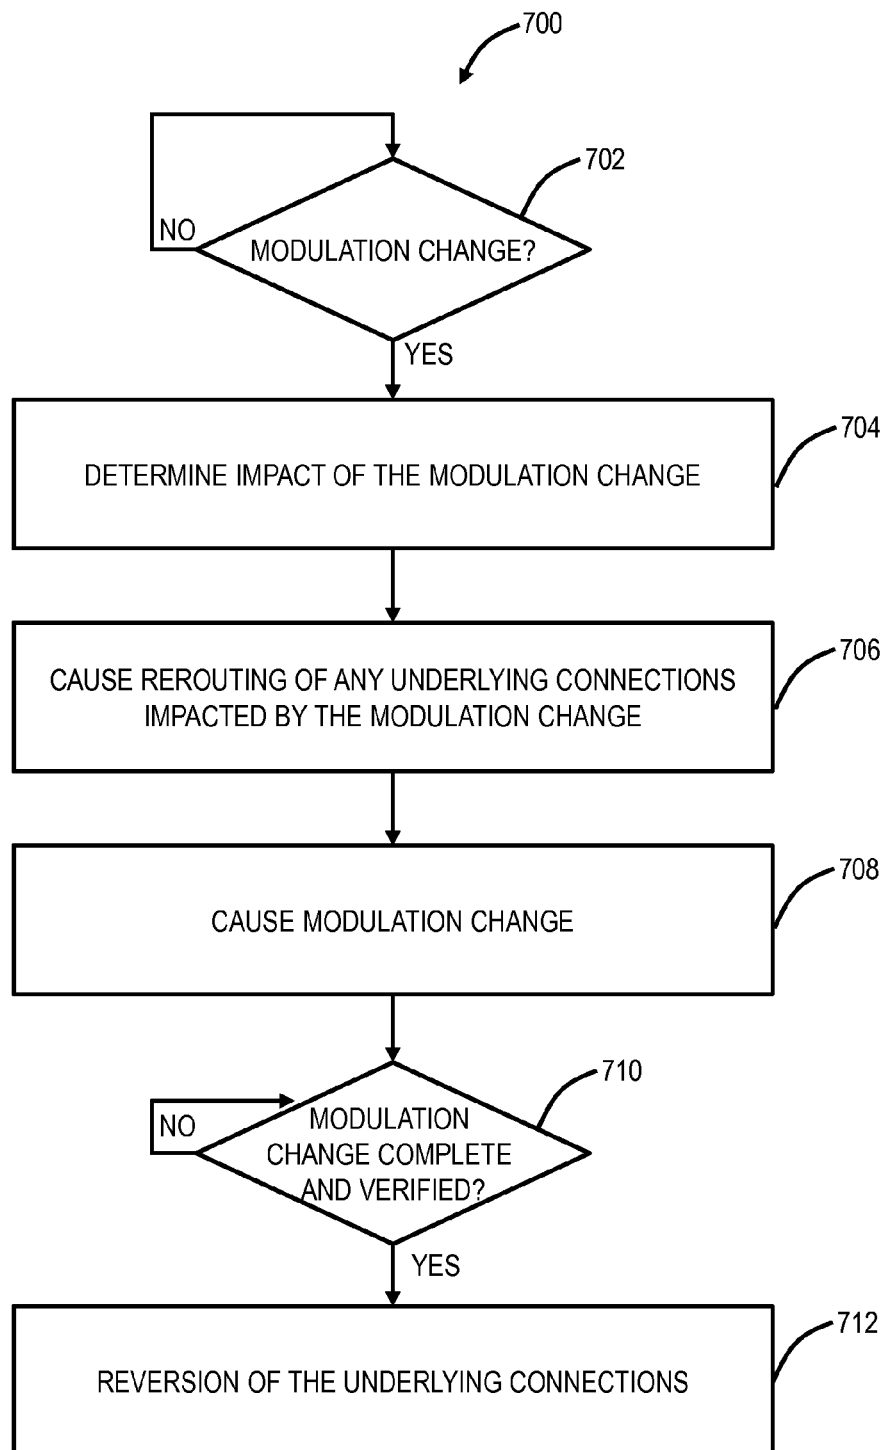


FIG. 8

**FIG. 9**

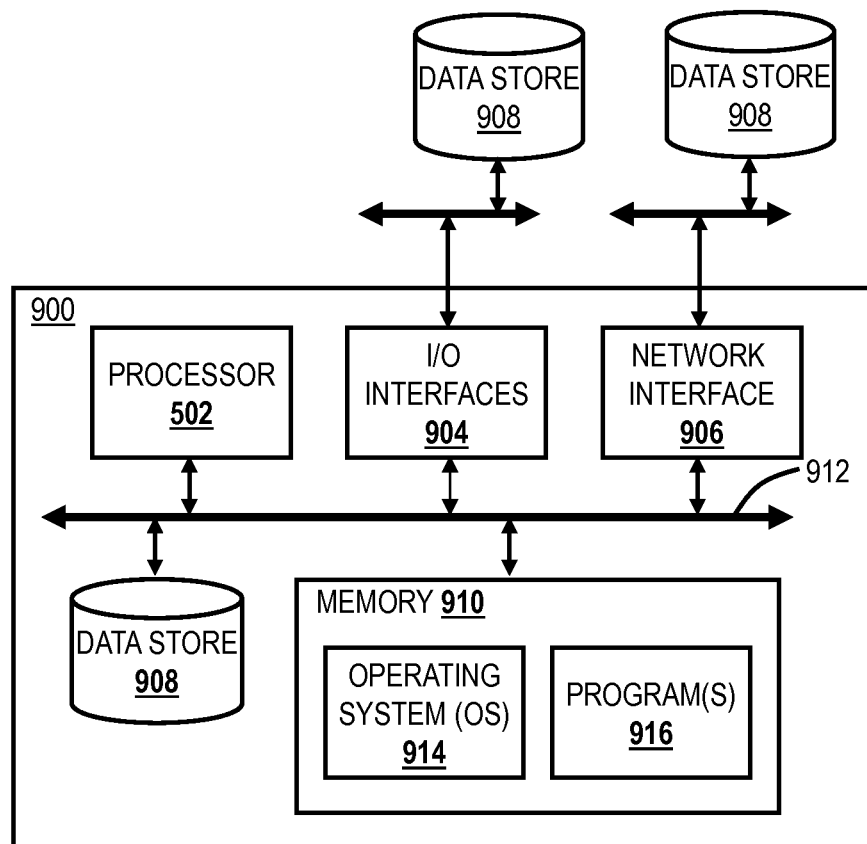


FIG. 10

1

HITLESS MODULATION SCHEME CHANGE SYSTEMS AND METHODS IN OPTICAL NETWORKS

CROSS-REFERENCE TO RELATED APPLICATION(S)

The present non-provisional patent application/patent is a continuation-in-part of U.S. patent application Ser. No. 14/176,908, filed on Feb. 10, 2014, and entitled "SYSTEMS AND METHODS FOR MANAGING EXCESS OPTICAL CAPACITY AND MARGIN IN OPTICAL NETWORKS," the contents of which are incorporated in full by reference herein.

FIELD OF THE DISCLOSURE

The present disclosure relates generally to optical networking. More particularly, the present disclosure relates to systems and methods for hitless modulation scheme changes in optical networks.

BACKGROUND OF THE DISCLOSURE

Fiber optic communication networks are experiencing rapidly increasing growth in capacity. This capacity growth is reflected by individual channel data rates, scaling from 10 Gbps (gigabits per second), to 40 Gbps, to developing 100 Gbps, and to future projections of 1000 Gbps channels and beyond. The capacity growth is also reflected by increasing total channel count and/or optical spectrum carried within an optical fiber. In the past, optical channels were deployed with a fixed capacity in terms of bandwidth as well as a fixed amount of overhead for forward error correction (FEC). For example, in a conventional system deployment, channels are deployed at 10 Gbps or 40 Gbps (plus associated overhead for FEC). These channels are designed to provide fixed data throughput capacity of 10 Gbps or 40 Gbps. Moreover, the performance limits of these channels are established assuming that the system is operating at full capacity, with all the optical channels present. The first in channels will operate in much more benign condition and have significant extra margin available. This margin is not utilized until much later in the life cycle of the system. For example, a single wavelength deployed on a new optical line system could have more than 10 dB of excess margin that is not currently utilized (without adding new hardware). This un-used margin can be considered wasted and forcing the system to operate in a non-cost effective way. If this extra margin could be utilized, even in a temporary way, to enhance data throughput of the modem for example, the economics of the system would be significantly improved.

Of note, next generation optical modems are equipped with the capability to support variable data throughput applications. Moreover, this capability will be provisionable. Therefore, depending on the opportunity, it would be advantageous to provision a modem at a higher data throughput when extra margin is available on new and low channel count deployments, usage of these next generation modem will allow to mine and use this excess margin and wasted capacity without requiring additional hardware. However, this excess margin will disappear as the channel counts approach full fill. It would be advantageous to have systems and methods for managing excess optical capacity and margin in optical networks in view of the above.

Fiber optic communication networks today are pushing up against the Shannon Limit within the non-linear tolerance of

2

the transponder technology currently in use. There is great interest in providing the best spectral efficiency possible, which is leading to the development of adaptive modulation techniques applied to fiber optic transmission. In wireless and Digital Subscriber Loop (DSL) technology, it is quite common to use adaptive modulation schemes which adapt to link conditions, e.g. High-Speed Downlink Packet Access (HSDPA) and Asymmetric digital subscriber line (ADSL2+). In optical, some latest generation transponders on the market are capable of changing modulation scheme, e.g. Ciena's WaveLogic3 family. Transponders in the future will be able to change modulation scheme more quickly, and may be optimized to do so. However, today's systems cannot take advantage of these in a hitless manner.

Although wireless and DSL technologies can react to channel conditions using adaptive modulation, the system constraints in fiber optic communication networks would not allow similar system techniques to work. In particular, there are two assumptions built into the algorithms used in these systems for wireless and DSL technologies. First, the data which is transported in both cases (wireless and DSL technologies) is bursty in nature, and the actual user data throughput vs. actual bit rate can be controlled. In wireless communication, the application layer is visible to the controller. In other words, the systems are designed to allow for periodic optimization where the resulting changes can be in modulation scheme. Optical transport networks are the core of the data network and as such see a rather continuous flow of traffic due to multiple levels of multiplexing and grooming. There are many sources of this data and the volume is also very high, so hold-offs on data transmission become too complex and/or expensive to implement.

Second, another simplifying condition is that the transmission time (distance) from modulator to demodulator is small compared to the baud rate of the transmission. In practical terms, for example, for HSDPA networks in a 5 km cell and the baud rate of 5 Mbps, there are 8.3 baud in flight at any time. In an optical network using 100 Gbps over 2000 km, there is 440 million baud in flight. This nearly 8 orders of magnitude difference represent a key difference in how to perform such changes. Although there are some transponder on the market today which can change the operational state to accommodate a different spectral efficiency which may be allowed by the link conditions, these cause some length of outage which can only be managed out of service or as a failure in the system, both of which have negative effects on the system, and drive operational complexity for the end user. Treating it as an outage may cause higher level protocols to attempt to recover from the failure, leaving the system vulnerable to further failures. The treatment as a failure may also cause re-transmission of data, etc. With the vast amount of data involved, this is simply unacceptable.

For example, using conventional optical modems, such as the WaveLogic3, testing was performed to switch in-service between Quadrature Phase Shift Keying (QPSK) and 16-Quadrature Amplitude Modulation (16-QAM). The switch requires several seconds because operating conditions on the line including non-linear impairments have to be calculated. Again, it is expected that the switch can be optimized, but likely not on the order of several milliseconds.

BRIEF SUMMARY OF THE DISCLOSURE

In an exemplary embodiment, a hitless modulation change method at a node in an optical network includes determining that a modulation change is warranted for an

optical modem in the node, the optical modem configured to communicate over an optical link; determining an impact of the modulation change on the optical link and associated underlying connections thereon; causing changes in a data plane for the associated underlying connections, prior to performing the modulation change; and causing the modulation change subsequent to accommodating the associated underlying connections in the data plane, thereby minimizing interruptions of the associated underlying connections due to the modulation change. The hitless modulation change method can further include causing reversion of the associated underlying connections to the optical link subsequent to completion and verification of the modulation change. The causing changes in the data plane can utilize a control plane and restoration of the associated underlying connections per normal control plane behavior. The normal control plane behavior can include any of mesh restoration via the control plane, 1+1/1:1 Automatic Protection Switching (APS), Subnetwork Connection Protection (SNCP), ring restoration, G.8032 Ethernet Ring Protection Switching (ERPS), and Virtual Local Area Network (VLAN) protection. The determining the impact can include determining a length of time the modulation change will take. The length of time the modulation change will take can be based on a type of the modulation change from first modulation scheme to a second modulation scheme, link conditions, and line measurements. The length of time can be communicated with the changes in the data plane for the associated underlying connections. The determining steps and the causing steps can be performed by a control agent communicatively coupled to the node. The control agent can operate in an autonomous manner and communicates with existing control plane functionality associated with the node to cause the changes in the data plane.

In another exemplary embodiment, a hitless modulation change system communicatively coupled to a node in an optical network includes a processor; and memory storing instructions that, when executed, cause the processor to determine that a modulation change is warranted for a optical modem in the node, the optical modem is configured to communicate over an optical link, determine an impact of the modulation change on the optical link and associated underlying connections thereon, cause changes in a data plane for the associated underlying connections, prior to performing the modulation change, and cause the modulation change subsequent to accommodating the associated underlying connections in the data plane, thereby minimizing interruptions of the associated underlying connections due to the modulation change. The memory storing instructions that, when executed, can further cause the processor to: cause reversion of the associated underlying connections to the optical link subsequent to completion and verification of the modulation change. The changes in the data plane can be caused utilizing a control plane and restoration of the associated underlying connections per normal control plane behavior. The normal control plane behavior can include any of mesh restoration via the control plane, 1+1/1:1 Automatic Protection Switching (APS), Subnetwork Connection Protection (SNCP), ring restoration, G.8032 Ethernet Ring Protection Switching (ERPS), and Virtual Local Area Network (VLAN) protection. The impact can include determining a length of time the modulation change will take. The length of time the modulation change will take can be based on a type of the modulation change from first modulation scheme to a second modulation scheme, link conditions, and line measurements. The length of time can be communicated with the changes in the data plane for the associated under-

lying connections. The hitless modulation change system can be implemented through or communicatively coupled to a Software Defined Networking (SDN) controller. The control agent can operate in an autonomous manner and communicates with existing control plane functionality associated with the node to cause the changes in the data plane.

In a further exemplary embodiment, an optical node implementing hitless modulation changes in an optical network includes one or more optical modems coupled to the optical network; a fabric coupled to the one or more optical modems for switching of connections; and a processing device implementing a control agent, wherein the control agent is configured to determine that a modulation change is warranted for an optical modem in the node, the optical modem is configured to communicate over an optical link, determine an impact of the modulation change on the optical link and associated underlying connections thereon, cause changes in a data plane for the associated underlying connections through the fabric, prior to performing the modulation change, and cause the modulation change subsequent to accommodating the associated underlying connections in the data plane, thereby minimizing interruptions of the associated underlying connections due to the modulation change. The changes can be caused in the data plane utilizing a control plane and restoration of the associated underlying connections per normal control plane behavior with the fabric.

BRIEF DESCRIPTION OF THE DRAWINGS

The present disclosure is illustrated and described herein with reference to the various drawings, in which like reference numbers are used to denote like system components/method steps, as appropriate, and in which:

FIG. 1 is a network diagram of an exemplary network for the systems and methods for managing excess optical capacity and margin in optical networks;

FIG. 2 are graphs of an example of spectral shaping fitting a 100G signal into 50 GHz of bandwidth and into 25 GHz of bandwidth or less;

FIG. 3 is a flowchart of a method for managing excess optical capacity and margin in optical networks;

FIG. 4 is a flowchart of another method for managing excess optical capacity and margin in optical networks;

FIG. 5 is a flowchart of a coexistence method for managing excess optical capacity and margin in optical networks with both variable capacity channels and fixed capacity channels intermixed;

FIG. 6 is a block diagram of an exemplary network element for use with the methods and systems described herein;

FIG. 7 is a block diagram of a controller to provide control plane processing and/or operations, administration, maintenance, and provisioning (OAM&P) for the network element of FIG. 6;

FIG. 8 is a block diagram of a hitless modulation change system at a node connected to an optical network; and

FIG. 9 is a flow chart of a hitless modulation change method which can be implemented through the control agent in coordination with the modem, the control plane, etc. from the hitless modulation change system of FIG. 8; and

FIG. 10 is a block diagram of a processing device that can operate the control agent from the hitless modulation change system of FIG. 8.

DETAILED DESCRIPTION OF THE DISCLOSURE

In various exemplary embodiments, hitless modulation scheme change systems and methods in optical networks are

described. The hitless modulation scheme change systems and methods allow modulation scheme changes of a transponder in a hitless manner so as to reap the benefit of changing the spectral efficiency to match link and system level requirements. The hitless modulation scheme change systems and methods provide a bottom-up approach and autonomous approach that negotiates a modulation change when one is warranted and ensures that existing traffic is not impacted during the modulation change. Due to the complex nature of high-speed optical transmission (e.g., coherent modulation, adaptive electrical signal processing, Forward Error Correction (FEC), etc.), it is extremely complex to switch a transponder or modem's modulation without impacting traffic. In this manner, the hitless modulation scheme change systems and methods propose an approach using a mechanism which coordinates between a data plane and control plane, i.e. using a fabric and a modem, to ensure current traffic is minimally impacted from a modulation scheme change. Note, as described herein, a hitless change means one in which traffic is impacted by 50 ms or less or by an amount based on service restoration (e.g., a mesh restoration event may be more than 50 ms). On the contrary, a modulation scheme change can take on the order of seconds to implement, given current implementations. It is expected that this amount of time will decrease, such as to 1 s or less, but still higher than 50 ms.

Additionally, in various exemplary embodiments, systems and methods for managing excess optical capacity and margin in optical networks are described. Fundamentally, the systems and methods exploit the fact that next-gen flexible optical modems can support various bit-rates well beyond a guaranteed bit-rate in most operating situations (i.e., the guaranteed bit-rate is engineered for full-fill or worst-case, and in all other situations, higher bit-rates typically can be achieved). In the systems and methods described herein, techniques are described to actively mine this excess capacity to provide additional bandwidth without additional hardware that can be used for various purposes such as restoration traffic, short-lived bandwidth-on-demand connections, or the like. In an exemplary aspect, the systems and methods described herein are advantageous in first-in builds in that this excess capacity can be used for restoration traffic without requiring additional hardware in low-fill deployments. This can significantly lower the costs of first-in builds. Further, the systems and methods described herein contemplate integration between the flexible optical modems; a management system, management plane, and/or control plane; a switching plane to enable use and management of this excess capacity as one or more logical interfaces.

Exemplary Network

Referring to FIG. 1, in an exemplary embodiment, a network diagram illustrates an exemplary network 10 implementing the systems and methods for managing excess optical capacity and margin in optical networks. The network 10 includes two interconnected network elements 12a, 12b via an optical link 14. Additionally, the optical link 14 can include additional components 16 which are omitted for illustration purposes. For example, the additional components 16 can include, without limitation, optical amplifiers, optical add/drop multiplexers (OADMs), reconfigurable OADMs (ROADMs), etc. In the context of the systems and methods, the network elements 12a, 12b are connected via the optical link 14 which is all-optical between the network elements 12a, 12b, i.e. no optical-electrical-optical (OEO) conversions between the network elements 12a, 12b. The optical link 14 can be a single span or multiple spans with

intermediate amplifiers. Those of ordinary skill in the art will recognize that the network 10 can include other network elements 12a, 12b forming various architectures, i.e. mesh, rings, linear, etc. The network 10 is presented as a single optical link (optionally with the components 16) for an illustration of the systems and methods.

The optical link 14 can include N channels (or wavelengths), denoted as λ_1 - λ_N . For example, the number N can be the maximum supported channels on the optical link 14. Additionally, the number N can be variable with respect to flexible grid channels (e.g., channels taking an arbitrary and variable amount of spectrum). For example, N can be 44 for 100 GHz channel spacing, 88 for 50 GHz channel spacing, or any combination in between to deliver between 36 and 88 wavelengths with flexible grid channels. Other embodiments are also contemplated. From a link engineering perspective, the optical link 14 is designed and implemented day one to support the N channels. However, greenfield installation or first-in builds (i.e., new) typically only include one or a couple of channels. Also, it can often take time to move from the couple of channels to a full complement of the N channels on the optical link 14. This can be referred to as a forecast tolerant modeling scheme where the optical link 14 is designed to support a full-fill that will eventually be realized, but is likely not present in first-in builds. Thus, from a system capacity perspective, the optical link 14 has unutilized margin and capacity in the first-in builds and where the optical link 14 has less than N channels deployed thereon.

In the context of the N channels, the N channels are either fixed capacity channels or variable capacity channels depending on associated hardware at the network elements 12a, 12b forming each of the N channels. In an exemplary embodiment, the optical line 14 can include one or more fixed capacity channels, one or more variable capacity channels, and/or a combination thereof. Fixed capacity channels are implemented through optical transceivers, transponders, muxponders (i.e., M:N combiners), etc. Here, the fixed capacity channels do not have an ability to vary the bandwidth, i.e. a 10 Gbps transponder with fixed capacity can only support 10 Gbps worth of traffic, etc. Fixed capacity channels may also include fixed channel spacing (e.g., 50/100 GHz) (i.e., fixed grid channels) and fixed FEC overhead. For a fixed capacity channel, if a channel has X dB excess margin, there is no way the fixed capacity channel can make use of this excess margin, i.e. the fixed capacity channel hardware is not configured to vary the bandwidth.

Variable capacity channels are implemented through flexible optical modems. In contrast to the fixed capacity channels, variable capacity channels typically include adaptable coherent modulation or non-coherent modulation, adaptive FEC schemes, and spectral shaping. A flexible optical modem can support a variable amount of bandwidth, e.g. from x Gbps to y Gbps, where $x < y$. For example, a flexible optical modem can support a guaranteed rate, e.g. 40G, 100G, 400G, 1T, etc. along with a higher supported rate, e.g. 40G->100G, 100G->200G, 400G->1T, etc. The flexible optical modem utilizes the adaptable coherent modulation, adaptive FEC schemes, and spectral shaping to support the variable amount of bandwidth. The limitations on the upper bound of the variable amount of bandwidth are based on i) what the optical link 14 can support, ii) backplane interfaces in the network element 12a, 12b with the flexible optical modem, and iii) adaptive modulation formats supported. An example of a flexible optical modem is the WaveLogic 3 from Ciena Corporation, the assignee of the present appli-

cation/patent. Also, note the flexible optical modem may also be referred to as a transceiver, transponder, muxponder, etc.

With respect to adaptive coherent modulation, the flexible optical modem can support various different baud rates through software-programmable modulation formats. The flexible optical modem can support programmable modulation, or constellations with both varying phase and/or amplitude. In an exemplary embodiment, the flexible optical modem can support multiple coherent modulation formats such as, for example, i) dual-channel, dual-polarization (DP) binary phase-shift keying (BPSK) for 100G at submarine distances, ii) DP quadrature phase-shift keying (QPSK) for 100G at ultra long haul distances, iii) 16-quadrature amplitude modulation (QAM) for 200G at metro to regional (600 km) distances), or iv) dual-channel 16QAM for 400G at metro to regional distances. Thus, in this exemplary embodiment, the same flexible optical modem hardware can support 100G to 400G. With associated digital signal processing (DSP) in the flexible optical modem hardware, moving from one modulation format to another is completely software-programmable. In another exemplary embodiment, the flexible optical modem can support N-QAM modulation formats with and without dual-channel and dual-polarization where N can even be a real number and not necessarily an integer. Here, the flexible optical modem can support non-standard speeds since N can be a real number as opposed to an integer, i.e. not just 100G, 200G, or 400G, but variable speeds, such as 130G, 270G, 560G, etc. Furthermore, with the DSP and software programming, the capacity of the flexible optical modem can be adjusted upwards or downwards in a hitless manner so as to not affect the guaranteed rate.

With respect to the adaptive FEC schemes, the flexible optical modem can support a new soft-decision forward error correction (soft FEC) algorithm. The soft FEC can be software-programmable to adjust for low latency demands versus capacity/performance demands. The soft FEC uses variable-rate FEC codes which can take up variable amounts of an overall signal, e.g. 20%, 16%, 10%, 7%, etc. As is known in the art, the stronger the FEC, the more margin in dB is provided. In this manner, the soft FEC provides another opportunity to mine the excess capacity on a variable capacity channel. For example, assume a variable capacity channel is deployed with 20% FEC overhead with margin of 10 dB. The FEC can be reduced, e.g., to 10% to reduce the margin and provide excess capacity for use. The strong FEC may not be needed until more channels are added to the optical link 14. An example of a soft-decision forward error correction algorithm is described in Ghoo et al., "Rate-Adaptive Coding for Optical Fiber Transmission Systems," IEEE JOURNAL OF LIGHTWAVE TECHNOLOGY, VOL. 29, NO. 2, Jan. 15, 2011, the contents of which are incorporated by reference herein. Note, the fixed capacity channel hardware may also implement FEC as well as a soft FEC. However, as described herein, the fixed capacity channel hardware is distinguishable from the variable capacity channel hardware in that it does not support an ability to mine the excess capacity. Rather, the fixed capacity channel hardware only supports a single guaranteed rate.

With respect to spectral shaping, the flexible optical modems can operate in both fixed- and flexible-grid environments. Referring to FIG. 2, in an exemplary embodiment, a spectral diagram illustrates an example of fitting a 100G signal into 50 GHz of bandwidth (graph 20 representing a QPSK 100G signal), into 25 GHz of bandwidth (graph 22 representing a 16QAM 100G signal), and into less than 25

GHz of bandwidth (graph 22a representing a spectrally shaped 16QAM 100G signal). Note, the 16QAM 100G which uses half the baud rate of the QPSK 100G. If one is on a fixed grid, there is no gain in spectral efficiency, e.g. both signals fit into a 50 GHz channel. If one is allowed to change the channel spacing flexibly, then the spectral efficiency can be doubled, e.g. two 16QAM 100G signals in 50 GHz spacing. For example, in a first-in build solely with flexible optical modems, it may be advantageous to use a flexible-grid and space each 100G signal in the minimal amount of bandwidth. However, in an existing fixed-grid, it may be required to fit the 100G into 50 GHz of bandwidth. Here, in an exemplary embodiment, the systems and methods propose to intentionally harm fixed capacity channels with excess, but unusable margin to allow the flexible optical modem to use the excess margin.

Various, it is an exemplary objective of the systems and methods to mine this unutilized margin and capacity to lower first-in network cost by allowing network operators to defer deploying excess capacity. Specifically, through the flexible optical modems, the systems and methods leverage the ability of the lines to provide the restoration bandwidth thereby deferring the deployment of additional optical interfaces as well as provide excess capacity that can be utilized for lower priority services, bandwidth-on-demand, etc. Specifically, first-in builds have significant excess margin, and with the emergence of flexible optical modems, it is an objective to provide and manage the excess margin to provide excess capacity without additional hardware or management constraints. That is, the flexible optical modems can significantly reduce initial costs by providing extra capacity that can be used for restoration, short-lived on-demand connections, or excess capacity with lower service-level agreements (SLAs). In conjunction with the foregoing, the systems and methods also include integration of this extra capacity with a management system, management plane, and/or control plane in the network 10 or other networks.

Managing Excess Optical Capacity and Margin

Referring to FIG. 3, in an exemplary embodiment, a flowchart illustrates a method 100 for managing excess optical capacity and margin in optical networks. The method 100 contemplates operation in the network 10 and other optical networks including flexible optical modems for variable capacity channels (and optionally with fixed capacity channels present as well.) The method 100 can be implemented on a single channel or wavelength of a flexible optical modem. The method 100 can be implemented on multiple channels concurrently or in series. For example, in series, each iteration of the method 100 may affect each subsequent iteration as the increased bandwidth of one channel may reduce the excess margin of the next. Performed concurrently, the method 100 may be based on a local determination of excess margin at each flexible optical modem without regard for collocated channels. The concurrent method could be independent (as stated) or in concert. Independent meaning that it is done per-channel without regard for other collocated channels, and in concert meaning that the margin of each channel is calculated taking into account the effect of the other collocated channels. This requires a "master" or nodal controller to amalgamate the channel information and perform the calculation.

The method 100 includes determining excess margin relative to a nominal guaranteed rate of a flexible optical modem (step 102). The nominal guaranteed rate can be the rate at which the flexible optical modem is configured to operate with a full-fill on the associated optical line. Also,

the nominal guaranteed rate can be the rate that is guaranteed through link engineering to work under any conceivable condition on the optical line such as full-fill. The excess margin (in dB) is the extra margin that the flexible optical modem presently sees given the current conditions on the optical line (e.g., channel count). That is, the excess margin is determined relative to margin needed to ensure performance at a nominal guaranteed rate. As stated herein, it is expected that on first-in deployments, the flexible optical modem may see significant margin given the engineering requirement to design for worst case (i.e., full-fill).

With the determined excess margin, the method **100** includes increasing capacity of the flexible optical modem to consume most or all of the excess margin (step **104**). Thus, the flexible optical modem supports a nominal guaranteed rate for guaranteed bandwidth and an excess rate for excess bandwidth where the excess rate minus the nominal guaranteed rate equals the excess capacity. Here, the method **100** can use all of the excess margin or most of it, leaving a small amount (e.g., 1 dB or less) for cushion to ensure the nominal guaranteed rate.

Next, the method **100** includes mapping the excess capacity to one or more logical interfaces (step **106**). The logical interfaces are typically 1:1 mapped with physical interfaces. Specifically, the logical interfaces are used by a management system, management plane, and/or control plane to map physical interfaces onto the optical line. Exemplary logical interfaces can be defined in terms of bandwidth such as, for example, 155 Mbps (Synchronous Transport Signal-level 1 (STS-1) or VC3), N×155 Mbps (N×STS-1), 1 Gbps (GbE), 2.5 Gbps (OC-48/STM-1, OTU1, ODU1), 10 Gbps (OC-192/STM-64, OTU2, ODU2, 10 GbE), 40 Gbps (OC-768/STM-256, OTU3, ODU3, 40 GbE), 100 Gbps (OTU4, ODU4, 100 GbE), variable capacity ODUflex, and the like. The logical interfaces can also be defined by signal type such as, for example, sub-network connections (SNCs), label switched paths (LSPs), 2F/4F BLSRs, 1+1/1:1 APS lines, UPSRs, VPSRs, 0:1 unprotected lines, etc. That is, the logical interfaces represent anything that allows the management system, management plane, and/or control plane to utilize the physical excess capacity from the method **100** in a network along with various switches.

The management system, management plane, and/or control plane are configured to recognize the excess capacity in terms of the associated logical interfaces and to allow physical hardware at the network elements **12** to support these extra logical interfaces. From a hardware perspective, the extra logical interfaces are formed on the optical line via the flexible optical modems in accordance with the method **100**. At the network elements **12** or collocated therewith, switches can be configured to process the extra logical interfaces through associated switching fabrics. Again, the management system, management plane, and/or control plane recognize these additional logical interfaces as extra traffic without requiring additional hardware (assuming the switching fabrics can support the additional capacity). Note, FIGS. 6-7 illustrate an exemplary network element **12** and associated control module for use with the systems and methods described herein.

In an exemplary embodiment, the method **100** includes flagging the bandwidth created in the step **106** on the one or more logical interfaces as excess capacity. For example, the flagging can include notifying the management system, management plane, and/or control plane that the one or more logical interfaces are excess capacity. The reason is to flag to the management system, management plane, and/or control plane is that this capacity can disappear and this needs

to be accounted for. The method **100** can be periodically reiterated for each flexible optical modem. For example, the method **100** can be reiterated at set intervals or based on an occurrence such as channel additions/deletions on the optical lines or margin changes/erosion on the optical lines. With each iteration of the method **100**, it is possible that the one or more logical interfaces could disappear or increase. For example, if channels are added to a line and the method **100** is rerun, the excess capacity could be decreased since the additional channels will likely reduce the excess margin. With a reduction in the excess capacity, some or all of the logical interfaces based thereon could disappear as the flexible optical modem scales back bandwidth or returns to the nominal guaranteed rate.

Accordingly, in an exemplary aspect, the method **100** contemplates using these logical interfaces based on the excess capacity for restoration bandwidth in new or low-fill optical networks as well as for bandwidth-on-demand, i.e. short-lived SNCs or LSPs, etc., and lower cost bandwidth with minimal SLA requirements. Specifically, in first-in builds, the method **100** can significantly reduce costs using the logical interfaces based on the excess capacity as mesh restoration SNCs or LSPs. This can defer the cost of additional optical interfaces to form unused capacity that is dedicated for restoration. Thus, in first-in builds, all optical hardware can be utilized for revenue generation.

Referring to FIG. 4, in an exemplary embodiment, a flow chart illustrates another method **200** for managing excess optical capacity and margin in optical networks. The method **200** is similar to the method **100** and provides additional details. Similarly, the method **200** contemplates operation in the network **10** and other optical networks, including flexible optical modems for variable capacity channels (and optionally with fixed capacity channels present as well.) The method **200** can be implemented on a single channel or wavelength of a flexible optical modem. The method **200** can be implemented on multiple channels concurrently or in series. For example, in series, each iteration of the method **200** may affect each subsequent iteration as the increased bandwidth of one channel may reduce the excess margin of the next. Performed concurrently, the method **200** may be based on a local determination of excess margin at each flexible optical modem without regard for collocated channels. The concurrent method could be independent (as stated) or in concert. Independent meaning that it is done per-channel without regard for other collocated channels, and in concert meaning that the margin of each channel is calculated taking into account the effect of the other collocated channels. This requires a "master" or nodal controller to amalgamate the channel information and perform the calculation.

The method **200** includes computing or providing a route for a network demand (step **202**). The network demand is a guaranteed amount of bandwidth needed in the network between two optical network elements **12**, e.g. 10G, 40G, 100G, etc. The method **200** can receive an explicit route or calculate a route using control plane techniques. Next, the method **200** includes determining path viability for the route and the network demand for an ideal bit-rate using a forecast tolerant modeling scheme (step **204**). For example, this functionality can be performed in a management system, an optical modeling system, etc., and this functionality includes determining the guaranteed wavelength capacity under worst-case conditions such as at full-fill, etc. That is, the forecast tolerant modeling scheme ensures the network demand can be serviced by the route regardless of future constraints. The step **204** could also optionally include a

11

wavelength assignment. The selection of wavelength could take into account the selection of wavelengths currently available (not in use).

Next, the method 200 includes determining path viability for the route and a maximum supported capacity on the existing network (step 206). The step 206 could also optionally include a wavelength assignment. The selection of wavelength in this case could differ from the step 204 in that it could select wavelengths which maximize the potential excess bandwidth. For example, it could choose to separate wavelengths from those already in service or to allocate a different spectral width to the channel being routed. The wavelength assignment in this step could change the wavelength previously chosen in step 204. Alternatively, this excess bandwidth aware wavelength assignment could be applied in step 204.

The step 204 looks at the worst case, whereas the step 206 looks at current conditions (i.e., right now without adding in margin for added channels or end-of-life operation). The step 204 determines the guaranteed wavelength capacity while the step 206 determines the current maximum wavelength capacity. It is the delta between these two scenarios that constitutes the excess margin and capacity opportunities with flexible optical modems. From a computation perspective, assign the value determined in the step 206 as Max and:

$$\text{Max_Engineered} = \text{Max} - \delta_{\text{margin}}$$

where Max_Engineered is the maximum currently supported bandwidth, Max is the result of the step 206 (i.e., the physical maximum bandwidth), and δ_{margin} is a small engineering margin simply to avoid a signal degrade threshold and this value can be 0 or a small amount such as <1 dB. The result of the step 204 can be denoted as Guaranteed, i.e. the guaranteed wavelength capacity. Accordingly:

$$\text{Excess} = \text{Max_Engineered} - \text{Guaranteed}$$

where Excess is the additional excess capacity currently supported that can be mined by the method 200 (or the method 100).

Next, the method 200 includes installing and/or activating a wavelength in the network at the Max_Engineered rate with a logical interface thereon supporting the guaranteed rate for the network demand and one or more logical interfaces providing the excess capacity (step 208). Here, the method 200, similar to the method 100, can provide these one or more logical interfaces from the excess capacity to a management plane and a switching plane for use thereof as restoration capacity, bandwidth-on-demand (BOD), short-lived services, etc. The method 200 can implement the various functionality described in the method 100 as well for implementing the one or more logical interfaces from the excess capacity.

The method 200 will operate with the logical interface supporting the guaranteed rate and with the one or more logical interfaces providing the excess capacity until a margin erosion, signal degradation, or other change (step 210). Again, it is expected at the client layer that the logical interface for the guaranteed capacity can be used for any service request, but specifically long lived traffic. On the other hand, the client layer could use the excess capacity for any service request, but it would be prudent to only use it for temporary traffic (e.g. restoration traffic, bandwidth-on-demand with a known termination date and time, etc.). In the method 200, if there is margin erosion or a signal degrade crossing (step 210), the flexible optical modem can drop the excess capacity and hitlessly revert back to the guaranteed bit-rate (step 212). In this way, the excess margin is now

12

used to make up for the margin erosion or the signal degrade crossing and not for the excess capacity. The excess capacity is lost, but the guaranteed capacity is protected from the margin erosion or the signal degrade crossing.

Once stability has been achieved for a set amount of time (e.g., 5 minutes, 2 hours, etc.) (step 214), the method 200 can include performing path viability for the route and a maximum supported capacity on the existing network (step 216). The step 216 is similar to the step 206. Once it is determined what excess margin exists after stability, the flexible optical modem can hitlessly increase its rate based on the maximum supported rate from the step 216 (and the guaranteed rate from the step 204). Also, if a path completely fails, then an alternate path is computed (per typical control plane behavior) and installed. The method 200 can operate as well on the new alternate path. For example, the method 200 can be implemented subsequent to a protection switch after stability is achieved.

Referring to FIG. 5, in an exemplary embodiment, a flow chart illustrates a coexistence method 300 for managing excess optical capacity and margin in optical networks with both variable capacity channels and fixed capacity channels intermixed. Specifically, the coexistence method 300 contemplates operation along with the methods 100, 200 in the network 10 and other optical networks, including flexible optical modems for variable capacity channels and with fixed capacity channels present as well. The coexistence method 300 begins with an assumption that any excess margin on a fixed capacity channel is unusable as discussed herein. The coexistence method 300 looks for opportunities to reduce this unusable excess margin to increase the excess margin on variable capacity channels for increased excess capacity according to the methods 100, 200.

The method 300 includes one or more variable capacity channels operating or planned on being operated on a same optical line or link as one or more fixed capacity channels (step 302). The method 300 can be implemented at various stages—in new systems where just a couple of channels are used all the way up to full-fill. The method 300 checks if there is an excess margin for any of the fixed capacity channels (step 304), and if not, the method 300 ends (step 306). If there is an excess margin on any of the fixed capacity channels (step 304), the method 300 utilizes various techniques to mine this excess margin for the benefit of the variable capacity channels (which in turn can implement the methods 100, 200 whereas the fixed capacity channels cannot).

The method 300 can include positioning or ensuring the variable capacity channels are located adjacent to fixed capacity channels on the optical spectrum (step 308). That is, it is advantageous for the method 300 to have fixed capacity channels adjacent to the variable capacity channels as opposed to separating these channels on the spectrum. The method 300 includes intentionally increasing performance of the variable capacity channels at the expense of the fixed capacity channels (with excess margin) to reduce the excess margin for the fixed capacity channels while concurrently increasing the excess margin for the variable capacity channel (step 310). In a way, it can be said that the method 300 intentionally harms the fixed capacity channels to remove the excess margin so it can be used by the variable capacity channels.

The method 300 contemplates various options for adjusting both the fixed capacity channels with excess margin and the variable capacity channels. For example, the fixed capacity channels could be transmitted at lower output powers to make these channels less intrusive to neighboring variable

13

capacity channels and therefore increase performance of the neighboring variable capacity channels. Further, the fixed capacity channels could be transmitted at a reduced baud rate and increased signal density to transmit in a format that takes more OSNR but uses less spectrum. Also, the variable capacity channel can intrude into the spectrum of the fixed capacity channel. For example, in FIG. 2, the variable capacity channel can extend 10 GHz into each of its neighbors to support 70 GHz of bandwidth versus 50 GHz thereby providing an additional margin for the variable capacity channel.

The method 300 can install or increase bit-rate of the variable capacity channels based on the increased performance and margin "stolen" from the fixed capacity channels and map this excess capacity to one or more logical interfaces such as described in the methods 100, 200 (step 312). In an exemplary embodiment, a method includes determining excess margin relative to margin needed to insure performance at a nominal guaranteed rate associated with a flexible optical modem configured to communicate over an optical link; causing the flexible optical modem to consume most or all of the excess margin, wherein the capacity increased above the nominal guaranteed rate includes excess capacity; and mapping the excess capacity to one or more logical interfaces for use by a management system, management plane, and/or control plane. The method can further include utilizing the one or more logical interfaces by the management system, management plane, and/or control plane as one of restoration bandwidth or short-lived bandwidth-on-demand connections. The method can further include determining the excess margin relative to the nominal guaranteed rate through the steps of: determining path viability of a network demand over the optical link for an ideal bit-rate using a forecast tolerant modeling scheme; determining path viability for a maximum supported capacity over the optical link based on existing conditions on the optical link; and determining the excess margin as a difference between the path viability for a maximum supported capacity and the path viability of the network demand along with including a small engineering margin. The method can further include detecting margin erosion or a signal degrade on the flexible optical modem; and dropping the excess capacity and hitlessly reverting to the nominal guaranteed rate. The method can further include after a period of stability subsequent to the margin erosion or the signal degrade, determining again the path viability for a new maximum supported capacity over the optical link based on existing conditions on the optical link; and hitlessly increasing a rate of the flexible optical modem based on the new maximum supported capacity.

The method can further include updating the determined excess margin relative to the nominal guaranteed rate in the flexible optical modem responsive to channels added or deleted on the optical link. The flexible optical modem can form a variable capacity channel, wherein the optical link can include a fixed capacity channel adjacent to the variable capacity channel, and the method can further include determining excess margin for the fixed capacity channel that is unusable since the fixed capacity channel cannot modify its rate; increasing performance of the variable capacity channel and/or decreasing performance of the fixed capacity channel based on the excess margin for the fixed capacity channel; and increasing bit-rate of the variable capacity channel based on margin gained by the increased performance of the variable capacity channel and/or the decreased performance of the fixed capacity channel. The method can further include increasing performance of the variable

14

capacity channel including extending associated optical spectrum into optical spectrum from the fixed capacity channel; and decreasing performance of the fixed capacity channel includes any of lowering output power to make the fixed capacity channel less intrusive to the variable capacity channel, or transmitting at a reduced baud rate and/or increased signal density to transmit in a format that uses less of the optical spectrum. The method can further include operating a control plane; and utilizing the one or more logical interfaces for restoration sub-network connections or label switched paths.

In another exemplary embodiment, a network element includes at least one flexible optical modem; and a controller configured to: determine excess margin relative to margin needed to insure performance at a nominal guaranteed rate associated with the at least one flexible optical modem configured to communicate over an optical link; cause the at least one flexible optical modem to consume most or all of the excess margin, wherein the capacity increased above the nominal guaranteed rate includes excess capacity; and map the excess capacity to one or more logical interfaces for use by a management system, management plane, and/or control plane. The controller can be further configured to utilize the one or more logical interfaces as one of restoration bandwidth or short-lived bandwidth-on-demand connections. The controller can be further configured to determine the excess margin relative to the nominal guaranteed rate through the steps of: determine path viability of a network demand over the optical link for an ideal bit-rate using a forecast tolerant modeling scheme; determine path viability for a maximum supported capacity over the optical link based on existing conditions on the optical link; and determine the excess margin as a difference between the path viability for a maximum supported capacity and the path viability of the network demand along with including a small engineering margin. The controller can be further configured to detect margin erosion or a signal degrade on the at least one flexible optical modem; and drop the excess capacity and hitlessly reverting to the nominal guaranteed rate.

The controller can be further configured to, after a period of stability subsequent to the margin erosion or the signal degrade, determine again the path viability for a new maximum supported capacity over the optical link based on existing conditions on the optical link; and hitlessly increase a rate of the flexible optical modem based on the new maximum supported capacity. The controller can be further configured to update the determined excess margin relative to the nominal guaranteed rate in the at least flexible optical modem responsive to channels added or deleted on the optical link. The at least one flexible optical modem can form a variable capacity channel, wherein the optical link includes a fixed capacity channel adjacent to the variable capacity channel, and the controller can be further configured to determine excess margin for the fixed capacity channel that is unusable since the fixed capacity channel cannot modify its rate; increase performance of the variable capacity channel and/or decreasing performance of the fixed capacity channel based on the excess margin for the fixed capacity channel; and increase a bit-rate of the variable capacity channel based on margin gained by the increase performance of the variable capacity channel and/or the decreased performance of the fixed capacity channel. The controller can be further configured to increase performance of the variable capacity channel including extending associated optical spectrum into optical spectrum from the fixed capacity channel; and decrease performance of the fixed

15

capacity channel including the of lowering output power to make the fixed capacity channel less intrusive to the variable capacity channel, or transmitting at a reduced baud rate and/or increased signal density to transmit in a format that uses less of the optical spectrum. The controller can be

In yet another exemplary embodiment, a network includes a plurality of interconnected network elements, at least one link in the network formed between two of the plurality of interconnected network elements is formed by flexible optical modems; a control plane communicatively coupled to the plurality of interconnected network elements; and a controller communicatively coupled to the flexible optical modems and configured to: determine excess margin needed to insure performance at a nominal guaranteed rate over the at least one link; cause the flexible optical modems to consume most or all of the excess margin, wherein the capacity increased above the nominal guaranteed rate includes excess capacity; and map the excess capacity to one or more logical interfaces for use by the control plane. The controller can be further configured to utilize the one or more logical interfaces as one of restoration bandwidth or short-lived bandwidth-on-demand connections.

Exemplary Network Element

Referring to FIG. 6, in an exemplary embodiment, a block diagram illustrates an exemplary network element 12 for use with the methods and systems described herein. In an exemplary embodiment, the exemplary network element 12 can be a network element that may consolidate the functionality of a multi-service provisioning platform (MSPP), digital cross connect (DCS), Ethernet and/or Optical Transport Network (OTN) switch, dense wave division multiplexed (DWDM) platform, etc. into a single, high-capacity intelligent switching system providing Layer 0, 1, and 2 consolidation. In another exemplary embodiment, the network element 12 can be any of an OTN add/drop multiplexer (ADM), a SONET/SDH/OTN ADM, a multi-service provisioning platform (MSPP), a digital cross-connect (DCS), an optical cross-connect, an optical switch, a router, a switch, a wavelength division multiplexing (WDM) terminal, an access/aggregation device, etc. That is, the network element 12 can be any digital system with ingress and egress digital signals and switching therebetween of channels, timeslots, tributary units, wavelengths, etc. utilizing OTN, SONET, SDH, etc. Alternatively, the network element 12 can exclude digital switching and solely provide optical switching and/or transmission. While the network element 12 is generally shown as an optical network element, the systems and methods contemplated for use with any switching fabric, network element, or network based thereon.

In an exemplary embodiment, the network element 12 includes common equipment 410, one or more line modules 420, and one or more switch modules 430. The common equipment 410 can include power; a control module; operations, administration, maintenance, and provisioning (OAM&P) access; user interface ports; and the like. The common equipment 410 can connect to a management system 450 through a data communication network 460. The management system 450 can include a network management system (NMS), element management system (EMS), or the like. Additionally, the common equipment 410 can include a control plane processor configured to operate a control plane as described herein. The network element 12 can include an interface 470 for communicatively coupling the common equipment 410, the line modules 420, and the

16

switch modules 430 therebetween. For example, the interface 470 can be a backplane, mid-plane, a bus, optical or electrical connectors, or the like. The line modules 420 are configured to provide ingress and egress to the switch modules 430 and external to the network element 12. In an exemplary embodiment, the line modules 420 can form ingress and egress switches with the switch modules 430 as center stage switches for a three-stage switch, e.g. a three stage Clos switch. Other configurations and/or architectures are also contemplated. The line modules 420 can include optical transceivers, such as, for example, 1 Gbps (GbE PHY), 2.5 Gbps (OC-48/STM-1, OTU1, ODU1), 10 Gbps (OC-192/STM-64, OTU2, ODU2, 10 GbE PHY), 40 Gbps (OC-768/STM-256, OTU3, ODU3, 40 GbE PHY), 100 Gbps (OTU4, ODU4, 100 GbE PHY), etc.

Further, the line modules 420 can include a plurality of optical connections per module and each module may include a flexible rate support for any type of connection, such as, for example, 155 Mbps, 622 Mbps, 1 Gbps, 2.5 Gbps, 10 Gbps, 40 Gbps, 100 Gbps, 400 Gbps, 1 Tbps, and any rate in between. The line modules 420 can include wavelength division multiplexing interfaces, short reach interfaces, and the like, and can connect to other line modules 420 on remote network elements, end clients, edge routers, and the like. From a logical perspective, the line modules 420 provide ingress and egress ports to the network element 12, and each line module 420 can include one or more physical ports. As described herein the line modules 420 can support either fixed capacity channels or variable capacity channels. The line modules 420 can be transponders, muxponders, flexible optical modems, etc. Note, if the network element 12 is a DWDM terminal, the switch modules 430 may be omitted and the line modules 420 act as transponders, muxponders, etc. It is assumed that a switch device is at some point connected to the DWDM terminal to support the one or more logical interfaces that are formed from the excess capacity.

The switch modules 430 are configured to switch channels, timeslots, tributary units, etc. between the line modules 420. For example, the switch modules 430 can provide wavelength granularity (Layer 0 switching), SONET/SDH granularity such as Synchronous Transport Signal-1 (STS-1) and variants/concatenations thereof (STS-n/STS-nc), Synchronous Transport Module level 1 (STM-1) and variants/concatenations thereof, Virtual Container 3 (VC3), etc.; OTN granularity such as Optical Channel Data Unit-1 (ODU1), Optical Channel Data Unit-2 (ODU2), Optical Channel Data Unit-3 (ODU3), Optical Channel Data Unit-4 (ODU4), Optical Channel Data Unit-flex (ODUflex), Optical channel Payload Virtual Containers (OPVCs), ODTUGs, etc.; Ethernet packet granularity; Digital Signal n (DSn) granularity such as DS0, DS1, DS3, etc.; and the like. Specifically, the switch modules 430 can include both Time Division Multiplexed (TDM) (i.e., circuit switching) and packet switching engines. The switch modules 430 can include redundancy as well, such as 1:1, 1:N, etc. In an exemplary embodiment, the switch modules 430 provide OTN, SONET, or SDH switching.

Those of ordinary skill in the art will recognize the network element 12 can include other components which are omitted for illustration purposes, and that the systems and methods described herein are contemplated for use with a plurality of different network elements with the network element 12 presented as an exemplary type of network element. For example, in another exemplary embodiment, the network element 12 may not include the switch modules 430, but rather have the corresponding functionality in the

line modules **420** (or some equivalent) in a distributed fashion or completely omit the corresponding functionality as in the case of a DWDM terminal. For the network element **12**, other architectures providing ingress, egress, and switching therebetween are also contemplated for the systems and methods described herein. In general, the systems and methods described herein contemplate use with any network element providing switching and/or transport of OTN, SONET, SDH, etc. channels, timeslots, tributary units, wavelengths, packets, etc. Furthermore, the network element **12** is merely presented as one exemplary implementation for the systems and methods described herein. Those of ordinary skill in the art will recognize the systems and methods can be used for practically any type of network element that includes flexible optical modems for supporting variable capacity channels.

Exemplary Controller

Referring to FIG. 7, in an exemplary embodiment, a block diagram illustrates a controller **500** to provide control plane processing and/or operations, administration, maintenance, and provisioning (OAM&P) for the network element **12**. The controller **500** can be part of common equipment, such as common equipment **410** in the network element **12**. The controller **500** can include a processor **502** which is hardware device for executing software instructions such as operating the control plane. The processor **502** can be any custom made or commercially available processor, a central processing unit (CPU), an auxiliary processor among several processors associated with the controller **500**, a semiconductor-based microprocessor (in the form of a microchip or chip set), or generally any device for executing software instructions. When the controller **500** is in operation, the processor **502** is configured to execute software stored within memory, to communicate data to and from the memory, and to generally control operations of the controller **500** pursuant to the software instructions. The controller **500** can also include a network interface **504**, a data store **506**, memory **508**, an I/O interface **510**, and the like, all of which are communicatively coupled therebetween and with the processor **502**.

The network interface **504** can be used to enable the controller **500** to communicate on a network, such as to communicate control plane information to other controllers, to the management system **460**, to a Software Defined Networking or OpenFlow controller, and the like. The network interface **504** can include, for example, an Ethernet card (e.g., 10BaseT, Fast Ethernet, Gigabit Ethernet) or a wireless local area network (WLAN) card (e.g., 802.11a/b/g). The network interface **504** can include address, control, and/or data connections to enable appropriate communications on the network. The data store **506** can be used to store data, such as control plane information, provisioning data, OAM&P data, etc. The data store **506** can include any of volatile memory elements (e.g., random access memory (RAM, such as DRAM, SRAM, SDRAM, and the like)), nonvolatile memory elements (e.g., ROM, hard drive, flash drive, CDROM, and the like), and combinations thereof. Moreover, the data store **506** can incorporate electronic, magnetic, optical, and/or other types of storage media. The memory **508** can include any of volatile memory elements (e.g., random access memory (RAM, such as DRAM, SRAM, SDRAM, etc.)), nonvolatile memory elements (e.g., ROM, hard drive, flash drive, CDROM, etc.), and combinations thereof. Moreover, the memory **508** may incorporate electronic, magnetic, optical, and/or other types of storage media. Note that the memory **508** can have a distributed

architecture, where various components are situated remotely from one another, but may be accessed by the processor **502**.

The I/O interface **510** includes components for the controller **500** to communicate to other devices in a node, such as through the local interface **514**. The components (**502**, **504**, **506**, **508**, **510**) are communicatively coupled via a local interface **514**. The local interface **514** and the I/O interface **510** can be, for example but not limited to, one or more buses or other wired or wireless connections, as is known in the art. The local interface **514** and the I/O interface **510** can have additional elements, which are omitted for simplicity, such as controllers, buffers (caches), drivers, repeaters, and receivers, among many others, to enable communications. Further, the local interface **514** and the I/O interface **510** can include address, control, and/or data connections to enable appropriate communications among the aforementioned components.

The network element **12**, the controller **500**, and associated optical networks and the like can utilize control plane systems and methods in addition to or in place of the standard management system functionality. Control plane systems and methods provide automatic allocation of network resources in an end-to-end manner. Exemplary control planes may include Automatically Switched Optical Network (ASON) as defined in G.8080/Y.1304, Architecture for the automatically switched optical network (ASON) (February 2005), the contents of which are herein incorporated by reference; Generalized Multi-Protocol Label Switching (GMPLS) Architecture as defined in Request for Comments (RFC): 3945 (October 2004) and the like, the contents of which are herein incorporated by reference; Optical Signaling and Routing Protocol (OSRP) from Ciena Corporation which is an optical signaling and routing protocol similar to PNNI (Private Network-to-Network Interface) and MPLS; or any other type control plane for controlling network elements at multiple layers, and establishing connections there between. It is essential for the operation of control planes to have control plane signaling and Operations, Administration, Maintenance, and Provisioning (OAM&P) connectivity between nodes.

It is important to note that the one or more logical interfaces from the methods **100**, **200**, **300** are provided to the control plane and/or management system for use thereof. In the control plane, the one or more logical interfaces can be used for path computation, especially for mesh reroutes in failure scenarios or for short-lived SNC or LSP requests. Thus, the systems and methods described herein provide the methods **100**, **200**, **300** for determining and setting maximum bandwidths on flexible optical modems, the control plane and/or management system are made aware of the excess capacity as one or more logical interfaces that are flagged as excess capacity, and collocated switching fabrics can use these one or more logical interfaces at the direction of the control plane and/or management system for various purposes such as restoration or short-lived SNCs or LSPs. Hitless Modulation Change System

Referring to FIG. 8, in an exemplary embodiment, a block diagram illustrates a hitless modulation change system **600** at a node **12** connected to an optical network **602**. The node **12** includes one or more modems **610**, a fabric **620**, and, optionally, one or more client transceivers **630**. The modems **610** are flexible optical modems as described herein, and interface optically via a bi-directional connection **635** to the optical network **602**. The fabric **620** can include an optical switching fabric (wavelengths), a time division multiplexing (TDM) fabric (OTN, SONET, SDH, etc.), and/or a packet

switching fabric (Ethernet, MPLS, IP, etc.). Note, the fabric 620 can be implemented in various different devices for different layers and/or in an integrated manner. Here, in FIG. 8, the fabric 620 is shown as a single logical device for illustration purposes. The fabric 620 is configured to switch wavelengths, channels, timeslots, tributary units, and/or packets between and among the modems 610 as well as the client transceivers 630 (note, in some exemplary embodiments, the client transceivers 630 may be omitted in the node 12). The client transceivers 630 are for local add/drop of connections from the node 12.

The node 12 and the optical network 602 can include a control plane 640. Note, the connections formed through the fabric 620 between and among the modems 610 and/or the client transceivers 630 can be referred to as a data plane. The control plane 640 is a separate management plane to control the connections in the data plane. Again, as described herein, the control plane 640 can include ASON, GMPLS, OSRP, etc., as well as other management mechanisms such as Software Defined Networking (SDN), Path Computation Element (PCE), and the like. That is, the control plane 640 represents any mechanism for automatically allocating the connections within the optical network 602 through the node 12.

As described herein, the node 12 and the modems 610, as flexible optical modems, support variable and flexible modulation schemes meant to maximize optical capacity. For the hitless modulation scheme change systems and methods, the node 12 includes or is communicatively coupled to a control agent 650 with the ability to probe the optical network 602 to determine not only that a change is warranted and acceptable, but also to determine the parameters of the impact of the traffic carrying capability of the channels, and provide a mechanism to perform the change without impacting traffic. Specifically, the control agent 650 can be an autonomous system that coordinates between the modems 610 and the fabric 620 to ensure that any modulation change in a particular modem 610 is hitless to all underlying connections currently on the modem 610. This can be done through rerouting of the underlying connections via the fabric 620 and/or the control plane 640 during a modulation change process, and reverting the underlying connections back to the modem 610 subsequent to completion of the modulation change process. In this manner, the modulation change process is hitless to the underlying connections, the underlying connections are rerouted per existing control plane behavior for restorations, and an end user does not have to manually manage the modulation change process.

The control agent 650 contemplates operation locally or remotely relative to the node 12. In an exemplary embodiment, the control agent 650 is implemented in the controller 500. In another exemplary embodiment, the control agent 650 is implemented in the management system 450 or an SDN controller or application. Various other embodiments are also contemplated. The control agent 650 is communicatively coupled to the modems 610, to the control plane 640, the SDN controller, the management system 450, and/or to the fabric 620. The control agent 650 can be configured to perform various functions described herein, such as determining a modulation change is warranted and supported, determining the impact of the modulation change, causing rerouting of connections impacted by the modulation change, and the like. That is, the control agent 650 implements an autonomous self-optimization of the modems 610 in the optical network 602. Importantly, the control agent 650 works with existing functionality in the

optical network 602 and the control plane 640, treating the underlying connections as per standard restoration procedures while a modulation change occurs. In this manner, network behavior is preserved and the underlying connections are minimally impacted due to the modulation change (i.e., hitless based on the underlying restoration technique).

The approach contemplated herein can be referred to as a bottom-up approach in that a modulation change is determined, and the control agent 650 performs rerouting of the underlying connections first, with the control plane 640, etc., prior to implementing the modulation change. Thus, the modulation change can be implemented safely without impacting the underlying connections, without adjustments at run-time, and without needing to minimize the length of the modulation change. The control agent 650 can be informed once the modulation change is complete and verified, and cause reversion of the underlying connections. The control agent 650 allows a bottom-up approach that is less complex (using existing functionality rather than adding extra functionality), is autonomous, and enables better temporal average spectral efficiency of the modems 610 and the optical network 602. Additionally, the control agent 650 can implement all or part of the methods 100, 200, 300 described herein.

Hitless Modulation Change Method

Referring to FIG. 9, in an exemplary embodiment, a flow chart illustrates a hitless modulation change method 700 which can be implemented through the control agent 650 in coordination with the modem 610, the control plane 640, etc. The hitless modulation change method 700 includes determining that a modulation change is needed on one of the modems 610 (step 702). This determination can be as described herein, e.g. excess margin available, decrease in the margin, etc. Note, the hitless modulation change method 700 is described with reference to a single modem 610; however, the hitless modulation change method 700 can be implemented concurrently for multiple modems 610 as required. The modulation change includes going from a first modulation to a second modulation, and the either the first modulation or the second modulation can provide more capacity, i.e. the modulation change can be an increase or decrease.

The hitless modulation change method 700 includes determining the impact of the modulation change (step 704). The impact includes which underlying connections are affected by the modulation change, a length of time of the modulation change (either estimated or calculated), and the like. Here, the hitless modulation change method 700 can determine which connections will be affected by the modulation change and about how long the modulation change will require. For example, the modem 610 can report a length of time the modulation change will take, based on the first modulation and the second modulation, link conditions, line measurements, etc.

Next, the hitless modulation change method 700 can cause rerouting of any underlying connections which will be impacted by the modulation change (step 706). Note, the hitless modulation change method 700 can also support local buffering of data and retransmission subsequent to completion of the modulation change; however, this is likely complex and costly due to the amount of bandwidth required. The hitless modulation change method 700 can cause rerouting by alerting the control plane 640 of which connections are affected, to allow the control plane 640 to reroute the connections per existing restoration techniques in the optical network 602, and optionally, provide the length of time so the control plane 640 could implement an

automatic reversion. Thus, the hitless modulation change method 700 utilizes the control plane 640 to coordinate a data path change in the data plane to accommodate the modulation change without interruption to the data flow.

Note, the rerouting of the underlying connections is per normal restoration techniques which may include mesh restoration via the control plane, 1+1/1:1 Automatic Protection Switching (APS), Subnetwork Connection Protection (SNCP), rings (BLSR, UPSR, etc.), G.8032 Ethernet Ring Protection Switching (ERPS) (ITU-T Recommendation G.8032/Y.1344 (02/12) Ethernet ring protection switching, the contents of which are incorporated by reference herein), Virtual Local Area Network (VLAN) protection, and the like. That is, the control plane can use whatever techniques are currently available in the data plane for rerouting/restoring connections in event of a fault. Here, the modulation change can be effectively viewed as a future fault that is mitigated in advance, ensuring no traffic impact.

However, a planned change in a channel path, in this case triggered by the desire to change modulation formats, can be much less of a hit, in terms of time, than a restoration/protection event driven by a failure. This is because the failure must be detected before the switch/re-route can take place. By this time, it's already too late, so to speak—as data has already been lost. In the case of a maintenance event, such as in this use case, one can perform what is called a “bridge and roll” where the new path is brought up by duplicating the data across the new path. In this case, even a very simple switch can be used, in which case the data which is “lost” is only due to the differential delay in the two paths. It is also possible to adjust or align to the new path before switching in which case there is no hit whatsoever. Thus, as described herein, the modulation change is hitless or substantially hitless (traffic loss is only due to the differential delay in the two paths which is typically less than 50 ms).

Once all the underlying connections are accounted for, either by rerouting or buffering, the hitless modulation change method 700 includes causing the modulation change (step 708). Here, the control agent 650 can signal to the modem 610 that it can implement the modulation change. The hitless modulation change method 700 waits for the modulation change to complete and verification (step 710). Note, it is expected the modulation change will take from about 1 s to multiple seconds. If there is a failure in the modulation change, the modem 610 can revert back to the first modulation. If there is a success, the modem 610 will operate with the second modulation, and associated components in the modem 610 will adjust and normalize to the line conditions with the second modulation (dispersion, non-linear effects, etc.). Once the second modulation is working, the modem 610 can verify successful operation through various OAM&P parameters such as FEC.

With proper operation of the second modulation or reversion back to the first modulation, the hitless modulation change method 700 can include reversion of the underlying connections back to the modem 610 (step 712). Here, the hitless modulation change method 700 can cause reversion or the underlying connections can automatically revert such as based on the length of time plus a buffer time. Again, the entire hitless modulation change method 700 can be autonomous meaning it requires no user intervention and it operates with existing data plane restoration functionality. SDN Controller/Service for the Application

Referring to FIG. 10, in an exemplary embodiment, a block diagram illustrates a processing device 900 that can operate the control agent 650. The processing device 900

can be a digital computer that, in terms of hardware architecture, generally includes a processor 902, input/output (I/O) interfaces 904, a network interface 906, a data store 908, and memory 910. It should be appreciated by those of ordinary skill in the art that FIG. 10 depicts the processing device 900 in an oversimplified manner, and a practical embodiment may include additional components and suitably configured processing logic to support known or conventional operating features that are not described in detail herein. The components (902, 904, 906, 908, and 910) are communicatively coupled via a local interface 912. The local interface 912 can be, for example but not limited to, one or more buses or other wired or wireless connections, as is known in the art. The local interface 912 can have additional elements, which are omitted for simplicity, such as controllers, buffers (caches), drivers, repeaters, and receivers, among many others, to enable communications. Further, the local interface 912 can include address, control, and/or data connections to enable appropriate communications among the aforementioned components.

The processor 902 is a hardware device for executing software instructions. The processor 902 can be any custom made or commercially available processor, a central processing unit (CPU), an auxiliary processor among several processors associated with the processing device 900, a semiconductor-based microprocessor (in the form of a microchip or chip set), or generally any device for executing software instructions. When the processing device 900 is in operation, the processor 902 is configured to execute software stored within the memory 910, to communicate data to and from the memory 910, and to generally control operations of the processing device 900 pursuant to the software instructions. The I/O interfaces 904 can be used to receive user input from and/or for providing system output to one or more devices or components. User input can be provided via, for example, a keyboard, touch pad, and/or a mouse. System output can be provided via a display device and a printer (not shown). I/O interfaces 904 can include, for example, a serial port, a parallel port, a small computer system interface (SCSI), a serial ATA (SATA), a fibre channel, Infiniband, iSCSI, a PCI Express interface (PCI-x), an infrared (IR) interface, a radio frequency (RF) interface, and/or a universal serial bus (USB) interface.

The network interface 906 can be used to enable the processing device 900 to communicate on a network. The network interface 906 can include, for example, an Ethernet card or adapter (e.g., 10BaseT, Fast Ethernet, Gigabit Ethernet, 10 GbE) or a wireless local area network (WLAN) card or adapter (e.g., 802.11a/b/g/n). The network interface 906 can include address, control, and/or data connections to enable appropriate communications on the network. A data store 908 can be used to store data. The data store 908 can include any of volatile memory elements (e.g., random access memory (RAM, such as DRAM, SRAM, SDRAM, and the like)), nonvolatile memory elements (e.g., ROM, hard drive, tape, CDROM, and the like), and combinations thereof. Moreover, the data store 908 can incorporate electronic, magnetic, optical, and/or other types of storage media. In one example, the data store 908 can be located internal to the processing device 900 such as, for example, an internal hard drive connected to the local interface 912 in the processing device 900. Additionally in another embodiment, the data store 908 can be located external to the processing device 900 such as, for example, an external hard drive connected to the I/O interfaces 904 (e.g., SCSI or USB connection). In a further embodiment, the data store 908 can

be connected to the processing device **900** through a network, such as, for example, a network attached file server.

The memory **910** can include any of volatile memory elements (e.g., random access memory (RAM, such as DRAM, SRAM, SDRAM, etc.)), nonvolatile memory elements (e.g., ROM, hard drive, tape, CDROM, etc.), and combinations thereof. Moreover, the memory **910** can incorporate electronic, magnetic, optical, and/or other types of storage media. Note that the memory **910** can have a distributed architecture, where various components are situated remotely from one another, but can be accessed by the processor **902**. The software in memory **910** can include one or more software programs, each of which includes an ordered listing of executable instructions for implementing logical functions. The software in the memory **910** includes a suitable operating system (O/S) **914** and one or more programs **916**. The operating system **914** essentially controls the execution of other computer programs, such as the one or more programs **516**, and provides scheduling, input-output control, file and data management, memory management, and communication control and related services. The one or more programs **916** may be configured to implement the various processes, algorithms, methods, techniques, etc. described herein.

In an exemplary embodiment, processing device **900** can be configured to implement the control agent **650**. For example, the processing device **900** can be an SDN controller or another device operating an SDN application in communication with the SDN controller. The processing device **900** can also be a standalone device in communication with the various components in the hitless modulation change system **600**. The SDN controller can also include an Application Programming Interface (API) which allows additional applications to interface with the SDN controller for data associated with the optical network **602**. In an exemplary embodiment, the control agent **650** can be implemented on the processing device **900** (or on the processing device **900** operating as the SDN controller) and receive data through the API. The processing device **900** can also be a standalone device in communication with the various components in the hitless modulation change system **600**. Other configurations are also contemplated.

The processing device **900** can include a queue that performs the hitless modulation change method **700** on one or more of the modems **610**. Again, the hitless modulation change systems and methods contemplate operation of the hitless modulation change method **700** for one or more of the modems **610** simultaneously, as appropriate. For example, the modems **610** can be managed based on location in the optical network **602** along with appropriate rules. One exemplary rule may be that only one modem **610** at a time in a node **12** can perform the hitless modulation change method **700**, whereas modems **610** in different nodes **12** can perform concurrently. Also, of note, the hitless modulation change method **700** is performed at both ends of a connection simultaneously. The control agent **650** can be a local device that communicates with a peer control agent **650** to coordinate the modulation change. Alternatively, the control agent **650** can have a global view and perform both ends simultaneously, for an optical section.

It will be appreciated that some exemplary embodiments described herein may include one or more generic or specialized processors ("one or more processors") such as microprocessors, digital signal processors, customized processors, and field programmable gate arrays (FPGAs) and unique stored program instructions (including both software and firmware) that control the one or more processors to

implement, in conjunction with certain non-processor circuits, some, most, or all of the functions of the methods and/or systems described herein. Alternatively, some or all functions may be implemented by a state machine that has no stored program instructions, or in one or more application specific integrated circuits (ASICs), in which each function or some combinations of certain of the functions are implemented as custom logic. Of course, a combination of the aforementioned approaches may be used. Moreover, some exemplary embodiments may be implemented as a non-transitory computer-readable storage medium having computer readable code stored thereon for programming a computer, server, appliance, device, etc. each of which may include a processor to perform methods as described and claimed herein. Examples of such computer-readable storage mediums include, but are not limited to, a hard disk, an optical storage device, a magnetic storage device, a ROM (Read Only Memory), a PROM (Programmable Read Only Memory), an EPROM (Erasable Programmable Read Only Memory), an EEPROM (Electrically Erasable Programmable Read Only Memory), Flash memory, and the like. When stored in the non-transitory computer readable medium, software can include instructions executable by a processor that, in response to such execution, cause a processor or any other circuitry to perform a set of operations, steps, methods, processes, algorithms, etc.

Although the present disclosure has been illustrated and described herein with reference to preferred embodiments and specific examples thereof, it will be readily apparent to those of ordinary skill in the art that other embodiments and examples may perform similar functions and/or achieve like results. All such equivalent embodiments and examples are within the spirit and scope of the present disclosure, are contemplated thereby, and are intended to be covered by the following claims.

What is claimed is:

1. A hitless modulation change method at a node in an optical network, the hitless modulation change method comprising:

determining that a modulation change from a first modulation scheme to a second modulation scheme is warranted for an optical modem in the node, the optical modem configured to communicate over an optical link;

determining an impact of the modulation change on the optical link and associated underlying connections thereon comprising determining a length of time the modulation change will take based on a type of the modulation change and conditions on the optical link; causing changes in a data plane for the associated underlying connections, prior to performing the modulation change; and

causing the modulation change subsequent to accommodating the associated underlying connections in the data plane, thereby minimizing interruptions of the associated underlying connections due to the modulation change.

2. The hitless modulation change method of claim 1, further comprising:

causing reversion of the associated underlying connections to the optical link subsequent to completion and verification of the modulation change.

3. The hitless modulation change method of claim 1, wherein the causing changes in the data plane utilizes a control plane and restoration of the associated underlying connections per normal control plane behavior.

25

4. The hitless modulation change method of claim 3, wherein the normal control plane behavior comprises any of mesh restoration via the control plane, 1+1/1:1 Automatic Protection Switching (APS), Subnetwork Connection Protection (SNCP), ring restoration, G.8032 Ethernet Ring Protection Switching (ERPS), and Virtual Local Area Network (VLAN) protection.

5. The hitless modulation change method of claim 1, wherein the length of time is communicated with the changes in the data plane for the associated underlying connections.

6. The hitless modulation change method of claim 1, wherein the determining steps and the causing steps are performed by a control agent communicatively coupled to the node.

7. The hitless modulation change method of claim 6, wherein the control agent operates in an autonomous manner and communicates with existing control plane functionality associated with the node to cause the changes in the data plane.

8. A hitless modulation change system communicatively coupled to a node in an optical network, the hitless modulation change system comprising:

- a processor; and
- memory storing instructions that, when executed, cause the processor to
 - determine that a modulation change from a first modulation scheme to a second modulation scheme is warranted for an optical modem in the node, the optical modem is configured to communicate over an optical link,
 - determine an impact of the modulation change on the optical link and associated underlying connections thereon comprising determining a length of time the modulation change will take based on a type of the modulation change and conditions on the optical link,
 - cause changes in a data plane for the associated underlying connections, prior to performing the modulation change, and
 - cause the modulation change subsequent to accommodating the associated underlying connections in the data plane, thereby minimizing interruptions of the associated underlying connections due to the modulation change.

9. The hitless modulation change system of claim 8, wherein the memory storing instructions that, when executed, further cause the processor to:

- cause reversion of the associated underlying connections to the optical link subsequent to completion and verification of the modulation change.

10. The hitless modulation change system of claim 8, wherein the changes in the data plane are caused utilizing a control plane and restoration of the associated underlying connections per normal control plane behavior.

11. The hitless modulation change system of claim 10, wherein the normal control plane behavior comprises any of mesh restoration via the control plane, 1+1/1:1 Automatic Protection Switching (APS), Subnetwork Connection Protection (SNCP), ring restoration, G.8032 Ethernet Ring Protection Switching (ERPS), and Virtual Local Area Network (VLAN) protection.

26

12. The hitless modulation change system of claim 8, wherein the length of time is communicated with the changes in the data plane for the associated underlying connections.

13. The hitless modulation change system of claim 8, wherein the hitless modulation change system is implemented through or communicatively coupled to a Software Defined Networking (SDN) controller.

14. The hitless modulation change system of claim 8, wherein the hitless modulation change system operates in an autonomous manner and communicates with existing control plane functionality associated with the node to cause the changes in the data plane.

15. An optical node implementing hitless modulation changes in an optical network, the optical node comprising:

- one or more optical modems;
- a fabric coupled to the one or more optical modems for switching of connections; and
- a processing device implementing a control agent, wherein the control agent is configured to
 - determine that a modulation change from a first modulation scheme to a second modulation scheme is warranted for an optical modem in the node, the optical modem is configured to communicate over an optical link,
 - determine an impact of the modulation change on the optical link and associated underlying connections thereon comprising determining a length of time the modulation change will take based on a type of the modulation change and conditions on the optical link,
 - cause changes in a data plane for the associated underlying connections through the fabric, prior to performing the modulation change, and
 - cause the modulation change subsequent to accommodating the associated underlying connections in the data plane, thereby minimizing interruptions of the associated underlying connections due to the modulation change.

16. The optical node of claim 15, wherein the changes are caused in the data plane utilizing a control plane and restoration of the associated underlying connections per normal control plane behavior with the fabric.

17. The optical node of claim 15, wherein the control agent is configured to

- cause reversion of the associated underlying connections to the optical link subsequent to completion and verification of the modulation change.

18. The optical node of claim 17, wherein the normal control plane behavior comprises any of mesh restoration via the control plane, 1+1/1:1 Automatic Protection Switching (APS), Subnetwork Connection Protection (SNCP), ring restoration, G.8032 Ethernet Ring Protection Switching (ERPS), and Virtual Local Area Network (VLAN) protection.

19. The optical node of claim 15, wherein the length of time is communicated with the changes in the data plane for the associated underlying connections.

20. The optical node of claim 15, wherein the control agent operates in an autonomous manner and communicates with existing control plane functionality associated with the optical node to cause the changes in the data plane.

* * * * *